REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)*	09-03-2006
2. REPORT TYPE	Final Technical Report
3. DATES COVERED *(From - To)*	01-10-2002 to 30-09-2005
4. TITLE AND SUBTITLE	Managing the disruptiveness of interruptions: Final technical report
5a. CONTRACT NUMBER	
5b. GRANT NUMBER	N00014-03-1-0063
5c. PROGRAM ELEMENT NUMBER	
5d. PROJECT NUMBER	05PR00047-00
5e. TASK NUMBER	
5f. WORK UNIT NUMBER	
6. AUTHOR(S)	Altmann, Erik M.
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	Department of Psychology Michigan State University East Lansing, MI 48824
8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	
10. SPONSOR/MONITOR'S ACRONYM(S)	
11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT	Approved for public release: Distribution is unlimited
13. SUPPLEMENTARY NOTES	
14. ABSTRACT	
15. SUBJECT TERMS	

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Erik Altmann
U	U	U			19b. TELEPHONE NUMBER *(Include area code)* 517-353-4406

Managing the Disruptiveness of Interruptions: Final Technical Report

Grant N00014-03-1-0063, performance period 01-Oct-2002 to 30-Sep-2005

Erik Altmann, 9 March 2006

Abstract

Task interruptions are pervasive in shipboard and other task environments of interest to the Navy. This report summarizes three years of experimental work examining the cognitive processes involved in the warning interval immediately prior to an interruption (in the moments between an alert, like the phone ringing, and the interruption proper, like the conversation with the caller), and immediately after an interruption, when the operator has to resume the interrupted task. Behavioral evidence suggests that people prepare before an interruption to resume afterwards, but that even with such preparation, resuming the interrupted task takes several seconds, net of baseline performance, and produces confusion in terms of reconstructing mental state associated with the interrupted task. Interventions to reduce the time cost of resumption have not been particularly successful, suggesting that interruptions are associated with fixed cognitive costs. Work continues on a separate grant to test this hypothesis with modeling and empirical methods.

Year 1 progress report

Interruptions — for example, to process communications from another operator or from automation — are prevalent in dynamic task environments, and this prevalence will only increase as fewer operators are asked to control more complex systems. As necessary as they are, however, interruptions can also produce confusion and disruption. For example, after answering a phone call, one may need to reconstruct where one was in the interrupted task, so as to resume without either skipping or redoing steps. The research program to be reported on examines cognitive and environmental variables affecting efficiency of resuming a complex task after an interruption (Trafton, Altmann, Brock, & Mintz, 2003). In the work to be described here, the focus is on the *interruption lag*, the brief period that often intervenes between an alert signaling a pending interruption, and the start of the interruption itself. (For an incoming phone call, the interruption lag would begin with the first ring and end with the start of the conversation with the caller.) The memory model on which the current work is based (Altmann & Trafton, 2002) makes strong predictions concerning the potential benefits of preparatory processing during the interruption lag. In particular, if resumption depends on remembering what one was doing prior to the interruption, the success of this remembering will depend critically on retrieval cues available both during the interruption lag (to allow encoding with cognitive state) and again at resumption time (to trigger retrieval of that state).

A series of four studies will be discussed that shows the benefits of having the full contents of the primary-task display available during the interruption lag, as compared with a situation in which the task display goes blank during the interruption lag. With interruption lags of six and eight seconds, display availability during the interruption lag significantly reduced resumption time (by roughly a second), whereas with interruption lags of two and four seconds, display availability during the interruption lag had no effect on resumption time. The results indicate that with several seconds available between an alert and the start of the interruption proper, people are able to find and somehow make use of environmental cues that help them pick up more efficiently where they left off.

These studies raise important questions about the nature of the cues that participants are focusing on and the nature of the processing into which these cues are incorporated. Studies currently in progress are aimed at addressing these questions through analysis of verbal and video protocol data. Interim goals are to develop more direct measures than simple response time of whether cue availability during the interruption lag improves resumption of suspended goals, and to identify simple heuristics or procedures suitable for developing into training interventions to reduce the interruption lag needed to make use of such cues.

References

Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science, 26*, 39-83.

Trafton, J. G., Altmann, E. M., Brock, D. P., & Mintz, F. E. (2003). Preparing to resume an interrupted tåsk: Effects of prospective goal encoding and retrospective rehearsal. *International Journal of Human-Computer Studies, 58*, 583-603.

Year 2 progress report

Year 2 of the project (June 2003 to May 2004) continued to focus on the warning interval preceding an interruption, and the effect of preparatory processes during this interval on time to resume the primary task after the interruption. Previous theoretical work on memory, cues, and forgetting had suggested that a warning interval could be used to prepare to resume the interrupted task after the interruption, much as people write "post-it" notes to themselves to remember something later. Studies during year 1 (October 2002 to May 2003) offered preliminary support for this hypothesis. These studies involved a cue condition, in which the primary-task display was preserved during the warning interval, and a nocue condition in which the primary task display was erased during the warning interval. With a sufficient warning interval (6 to 8 sec), the cue condition resumed the primary task about a second (20%) faster than the nocue condition. This indicated that some process, facilitated by cue availability, operated during the warning interval before the interruption to affect task resumption after the interruption, 30 to 40 sec later. These and related results will be presented at APS (Altmann & Trafton, 2004a) and have been submitted for presentation at Cognitive Science (Altmann & Trafton, 2004b).

In year 2, two tracks of empirical work were undertaken to explore the nature of preparation during the warning interval, one track focusing on basic questions about the nature of the underlying cognitive processes, and the other focusing on applied questions about how such preparation could be facilitated. On the basic track, a protocol study was conducted to acquire qualitative data on working-memory contents before and after an interruption. Fifty participants were involved, producing a database of think-aloud and mouse-movement protocols on 1500 interruptions; the data are currently being coded. A second study was undertaken to replicate the effect of cue availability, and to examine how it interacts with duration of the warning interval. The cue/nocue factor was crossed with 9 durations (0, 1, 2, ... 8 sec) in a between-subjects design. Data collection is projected to be complete by May 2004.

On the applied track, two studies were conducted to examine the effect of support for note-taking during the warning interval. The motivation was again the view that interruptions cause disruption by displacing primary-task information from working memory. In one study, participants in a record condition were asked to use the warning interval to record specific details about objects they were working on when the pending interruption was signaled; protocol data (from the study described above) had been examined to determine what information was thought about prevalently enough to warrant recording. Unexpectedly, the record condition resumed more slowly than no-record condition, suggesting that warning-interval preparation strategies are more diverse than can be supported by a single structured note-taking format. These results are described in a separate submission to the 2004 Cognitive Science conference (Clifford & Altmann, 2004; Jonathan Clifford is a student funded on this project). In a follow-up study, a free-form text notepad was made available as a strictly optional means for participants to record information they thought the interruption might cause them to forget. Notepad use occurred on 12% of interruptions, ranging from ranging from 0% to 70% across participants. These results offer some basic validation of the note-taking premise, but are consistent with the previous study in suggesting that underlying strategic and idiosyncratic variables are also at play. Future work will have to assess the perceived and actual benefits of notepad use, and examine whether training or other interventions can make them more uniform.

References

Altmann, E. M. & Trafton, J. G. (2004a). Task interruption: The cues we use to recover. Hot topic talk, 2004 annual convention of the American Psychological Society (Chicago, May 2004).

Altmann, E. M. & Trafton, J. G. (2004b). Task interruption and the role of cues. Paper submitted to the 2004 Annual Conference of the Cognitive Science Society.

Clifford, J. D. & Altmann, E. M. (2004). Managing multiple tasks: Reducing resumption time of the primary task. Paper submitted to the 2004 Annual Conference of the Cognitive Science Society.

Year 3 progress report

In previous years, we have shown a substantial time cost to resuming an interrupted task. This is computed by measuring resumption time, which is the interval between the moment the task environment for the interrupted task is reinstated and the moment of the operator's first subsequent action. This can be compared to per-action time, defined as the average interval separating two uninterrupted actions. If interruptions have no effect on response time, then resumption time and per-action time should be roughly equal. We have consistently found that resumption time is seconds longer than per-action time, indicating a substantial time cost to being interrupted.

What we have not previously addressed is the measurement of performance errors induced by interruptions. Error measures are an important complement to time measures. From an applied perspective, the effects of errors can range from inefficiency to catastrophe. From a theoretical perspective, error measures are necessary to interpret time measures; conceivably, the time costs of an interruption could reflect a temporary shift of emphasis away from speed and towards accuracy, instead of a generalized disruptive effect. However, error measures can be difficult to define in rich task environments, for example in ones like ours in which actions are reversible and the "quality" of an operator's performance is assessed by global, noisy means.

In the past year, we developed an error measure for performance in our task environment and showed that it converges with our time cost measure to indicate a generalized disruptive effect of interruption. After considerable experimentation, we found that redundant actions provided a reliable measure with good face validity. An action is redundant if it (a) is identical to the preceding action, and (b) does not change the state of the display or other aspect of the task environment. This measure can be automatically extracted from logfiles without subjective coding, and has face validity in that doing something over again, to no external effect, is objectively inefficient in terms of the number of actions taken. We have found that whereas the baseline frequency of redundant actions is about 0.5%, the frequency after interruptions is about 15%, which is a very large effect. Moreover, this redundancy measure is sensitive to at least one of the pre-interruption manipulations that we have also found to affect the time to resume the interrupted task. We anticipate that this high frequency of redundancies after an interruption will be an important source of constraint on models of post-interruption processing, because it suggests that, objective inefficiency aside, there is some underlying cognitive value, perhaps tied to retrieval of memory representations of state information tied to that action as a retrieval cue.

In other work, we have finished coding a large amount of verbal and mouse-movement protocol data on the moments before and after an interruption, in hopes of finding qualitative evidence of disruptive effects on the contents of working memory. We have also started developing a new task environment in the NASA-certified X-Plane flight simulator, with the aim of replicating and extending our findings in a task environment that is more dynamic and has improved ecological validity, and for which computational cognitive models exist that perform the task in place of human operators.